

Efficient Verification of Concurrent Programs Over the TSO Memory Model

Chinmay Narayan, Subodh Sharma, S.Arun-Kumar

Indian Institute of Technology Delhi

Abstract. We address the problem of efficient verification of multi-threaded programs running over Total Store Order (TSO) memory model. It has been shown that even with finite data domain programs, the complexity of control state reachability under TSO is non-primitive recursive. In this paper, we first present a bounded-buffer verification approach wherein a bound on the size of buffers is placed; verification is performed incrementally by increasing the size of the buffer with each iteration of the verification procedure until the said bound is reached. For programs operating on finite data domains, we also demonstrate the existence of a buffer bound k such that if the program is safe under that bound, then it is also safe for unbounded buffers. We have implemented this technique in a tool **ProofTraPar**. Our results against **memorax** [2], a state-of-the-art sound and complete verifier for TSO memory model, have been encouraging.

1 Introduction

The explosion in the number of schedules is central to the complexity of verifying the safety and correctness of concurrent programs. There exist a plethora of approaches in the literature that explore ways and means to address the schedule-space explosion problem; incidentally, many of the published techniques operate over the assumption of a sequentially consistent (SC) memory model. In contrast, almost all modern multi-core processors conform to memory models *weaker* than SC. A program executing on a relaxed memory model exhibits more behaviours than on the SC memory model. As a result, a program declared correct by a verification methodology that assumes SC memory model can possibly contain a buggy behaviour when executed on a relaxed memory model.

Consider x86 machines that conform to TSO (Total Store Ordering). The compiler or the runtime system of the program under the TSO memory model is allowed to reorder a read following a write (read and writes are to different variables) within a process, i.e. break the program order specified by the developer. Operationally, such a re-ordering is achieved by maintaining per-process store buffers. Write operations issued by a process/thread are enqueued in the store buffer local to that process. The buffered writes are later flushed (from the buffer) into the global memory. The point in time when flushes take place is deterministically known only when the store buffers are full. When the store

buffers are partially full, flushes are allowed to take place non-deterministically. Therefore, when a read operation of variable x , is executed by a process, the process first checks whether there is a recent write to x in the process's store buffer. If such a write exists then the value from store buffer is returned, otherwise the value is read from the global memory.

```

      flag1 = false, flag2 = false, t = 0;

      P1
      While(true){
1.  flag1:=true;
2.  t:=2;
3.  while(flag2 = true & t = 2);
4.  //Critical Section
5.  flag1:=false;
   }

      P2
      While(true){
6.  flag2:=true;
7.  t:=1;
8.  while(flag1 = true & t = 1);
9.  //Critical Section
10. flag2:=false;
   }

```

Fig. 1: Peterson's algorithm for two processes

Figure 1 shows Peterson's algorithm as an instance of a correct program under SC semantics but which can fail when executed under TSO. In this algorithm, two processes P_1 and P_2 coordinate their access to their respective critical sections using a shared variable t . This algorithm satisfies the mutual exclusion property under the SC memory model, i.e. both processes can not be

simultaneously present in their critical sections. The property however, is violated when the same algorithm is executed with a weaker memory model, such as TSO. Consider the following execution under TSO. The write operations at 1, 2, 6 and 7 from processes P_1 and P_2 are stored in store buffers and are yet to be reflected in the global memory. The reads at control locations 3 and 8 will return initial values, thereby violating the mutual exclusion property.

One can avoid such erroneous behaviors and restore the SC semantics of the program by inserting special instructions, called *memory fence*, at chosen control locations in the program. A memory fence ensures that the store buffer of the process (which executes the fence instruction) is flushed entirely before proceeding to the next instruction for execution. In the example, when *fence* instructions are placed after $\text{flag}_1 := \text{true}$ in P_1 and after $\text{flag}_2 := \text{true}$ in P_2 , the mutual exclusion property is restored.

Safety verification under TSO is a hard problem even in the case of finite data domain programs. The main reason for this complexity is the unboundedness of store buffers. A program can be proved correct under TSO only when the non-reachability of the error location is shown irrespective of the bound on the buffers. The work in [7] demonstrated the equivalence of the TSO-reachability problem to the coverability problem of lossy channel machines which is decidable and of non-primitive recursive complexity. A natural question is to ask if it is possible to have a buffer bound k such that if a finite data domain program is safe under the k -bounded TSO semantics then it is guaranteed to be safe even with unbounded buffers. For programs without loops such a statement seems to hold intuitively. For programs with loops, it is possible that a write instruction inside a loop keeps filling the buffer with values without ever getting them flushed to the main memory. However, for finite data domain programs, only a finite set

of different values will be present in this unbounded buffer and this leads to a sufficient bound on the buffer size.

In this paper we show that it is possible to verify a program P under TSO (with unbounded buffers) by generalizing the bounded buffer verification. Towards this we first define TSO_k , TSO semantics with buffer size k , and then characterize a bound k_0 such that if a program is safe in TSO_{k_0} then it is safe for any buffer bound greater than k_0 . We adapt a recently proposed trace partitioning based approach [16,25] for the TSO memory model. These methods work for the SC memory model as follows: the set of all SC executions of a program P are partitioned in a set of equivalence classes such that it is sufficient to prove the correctness of only one execution per equivalence class. As this trace partitioning approach works with symbolic executions, we first define an equivalent TSO semantics to generate a set of symbolic TSO traces. Subsequently we invoke a trace partitioning tool **ProofTraPar** [25] for proving the correctness of these traces. Note that the set of behaviors of a program P under TSO_k is a subset of the behaviors of P under $\text{TSO}_{k'}$ for any $k' > k$. The trace partitioning approach allows us to reuse the proof of correctness of P with buffer bound k in the proof of correctness of P with any buffer bound greater than k . In a nutshell, the main contributions of this work are:

- We characterize a buffer bound in case of finite state programs such that if the program is correct under TSO up to that bound then it is correct for unbounded buffers as well.
- We adapt the recently proposed trace partition based proof strategy of SC verification [16,25] for TSO by defining an equivalent TSO semantics to generate a set of symbolic TSO traces.
- We implement our approach in a tool, **ProofTraPar**[25], and compare its performance against **memorax**[2], a sound and complete verifier for safety properties under TSO. We perform competitively in terms of time as well as space. In a few examples, **memorax** timed out after consuming around 6GB of RAM whereas our approach could analyze the program in less than 100 MB memory.

Section 2 covers the related work in the area of verification under relaxed memory models. Section 3 covers the notations used in this paper. Section 4 shows the necessary and sufficient conditions to generalize bounded verification to unbounded buffers for finite data domain programs. Section 5 presents an equivalent TSO semantics to generate a set of symbolic traces which can be used by the trace partitioning tool **ProofTraPar** to check the correctness under a buffer bound. This section ends with an approach based on *critical cycle* to insert *memory fence* instructions. Section 6 compares the performance of our approach with **memorax**. Section 7 concludes with future directions.

2 Related Work

Figure 2 captures the related work in this area. Verification approaches for relaxed memory models can be broadly divided into three classes: precise, under-

approximate and over-approximate. For finite state programs, the work in [7,2] present sound and complete algorithms for control state reachability (finite state programs) under TSO and PSO memory models.

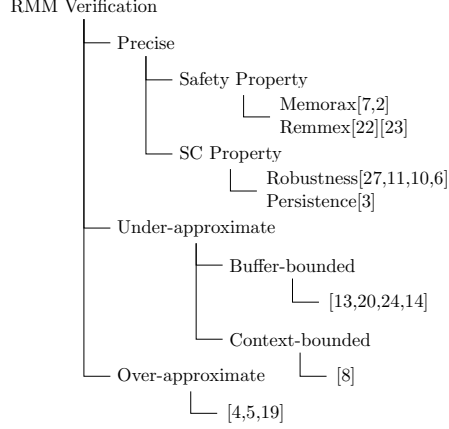


Fig. 2: State of the art

Sets of infinite configurations, arising from unbounded buffer size, are finitely presented using regular expressions. Acceleration based techniques that led to faster convergence in the presence of loops were presented in [22,23]. However, the termination of the algorithm was not guaranteed. Notice that in both [2] and [23] the specification was a set of control states to be avoided. One can also ask the state reachability question with respect to SC specification, i.e. does a program P reach only SC reachable states under a relaxed memory model? This problem was shown to be of the same complexity as of SC verification (Pspace-complete) and hence gave a more tractable cor-

rectness criterion than general state reachability problem. [27,11,10,6,3] work with this notion of correctness and give efficient algorithms to handle a range of memory models. In this paper we work with the control state reachability problem as opposed to the SC state reachability problem.

Over-approximate analyses [4,5,19] trade precision with efficiency and construct an over-approximate set of reachable states. Recently [1,28] used stateless model checking under TSO and PSO memory models. The main focus of these approaches are in finding bugs rather than proving programs correct. Another line of work to make the state reachability problem more tractable involved either restricting the size of buffers [13,20,24,14] or bounding the context switches [8] among threads. None of these methods give completeness guarantee even for the finite data domain programs.

3 Preliminary

A concurrent program is a set of processes uniquely identified by indices t from the set TID . As in [2,9], a process P_t is specified as an automaton $\langle Q_t, LABL_t, \delta_t, q_{0,t} \rangle$. Here Q_t is a finite set of control states, $\delta_t \subseteq Q_t \times LABL_t \times Q_t$ is a transition relation and $q_{0,t}$ is the initial state. Without loss of generality we assume every transition is labeled with a different symbol from $LABL_t$. $LABL_t$ represents a finite set of labels to symbolically represent the instructions of the program. Let SV be the set of shared variables of program P ranged over by x, y, z , Val be a finite set of constants ranged over by v , LV_t be the set of local variables of

process P_t ranged over by ℓ, m , and Exp_t be the set of expressions constructed using LV_t , Val and appropriate operators. Let $\text{LV} = \bigcup_t \text{LV}_t$, $\text{Exp} = \bigcup_t \text{Exp}_t$, and $\text{LABL} = \bigcup_t \text{LABL}_t$. Let a, b, c range over LABL and e range over Exp . Formally an instruction, from set INST , is one of the following type; (i) $x := e$, (ii) $\ell := x$, (iii) $\ell := e$, (iv) **assume**(e), and (v) **fence**, where $x \in \text{SV}$, $\ell \in \text{LV}$ and $e \in \text{Exp}$. A function $\text{Ins} : \text{LABL} \rightarrow \text{INST}$ assigns an instruction to every label.

The first two assignment instructions, (i) and (ii), are the write and the read operations of shared variables, respectively. Instruction (iii) assigns the value of an expression (constructed from local variables and constants) to a local variable, hence, does not include any shared memory operation. Instruction (iv) is used to model loop and conditional statements of the program. Note that the boolean expression e in **assume**(e) does not contain any shared variable. Instruction (v) represents the fence operation provided by the TSO architecture. Let $\text{Loc}(a)$ be the shared variable used in $\text{Ins}(a)$. For a function $F : A \times B$, let the function $F[p \leftarrow q]$ be the same as F everywhere except at p where it maps to q .

TSO Semantics In the TSO memory model, every process has a buffer of unbounded capacity. However, we present the TSO semantics by first defining a *k-bounded* TSO semantics where all buffers are of fixed size k . For a concurrent program P , the *k-bounded* semantics is given by a transition system $\text{TSO}_k = \langle S, \rightarrow_k, s_0 \rangle$. Every state $s \in S$ is of the form $(cs, \text{Lm}, \text{Gm}, \text{Buff}_k)$ where process control states $cs : \text{TID} \rightarrow Q$, $Q = \bigcup_t Q_t$, local memory $\text{Lm} : \text{TID} \times \text{LV} \rightarrow \text{Val}$, global memory $\text{Gm} : \text{SV} \rightarrow \text{Val}$, and *k-length* bounded buffers $\text{Buff}_k : \text{TID} \rightarrow (\text{SV} \times \text{Val})^k$. We overload operator ‘.’ to denote the concatenation of labels as well as a dereferencing operator to identify a specific field inside a state. Therefore, for a state s , $s.\text{Gm}$, $s.\text{Lm}$ and $s.\text{Buff}_k$ denote the functions representing global memory, local memory, and buffers respectively. Every write operation to a shared variable by process P_t initially gets stored in the process-local buffer provided that the buffer has less than k (buffer-bound) elements. This write operation is later removed from the buffer non-deterministically to update the global memory. A read operation of a shared variable say x , by a process P_t first checks the local buffer for any write to x . If buffer contains any write to x then the value of the last write to x is returned as a result of this read operation. If no such write is present in the buffer of P_t then the value is read from the global memory. A process executes instruction **fence** only when its local buffer is empty. For instruction **assume**(e), boolean expression e is evaluated in the local state of P_t . Execution proceeds only when the expression e evaluates to true. Assignment operation involving only local variables changes the local memory of P_t . The transition relation \rightarrow_k is defined in detail in the Appendix.

Relevance of the buffer size k Parameter k influences the extent of reordering that happens in an execution. For example, if $k = 0$ then no reordering happens and the set of executions is the same as under the SC memory model. Size parameter k , under the TSO memory model, allows any two instructions separated by at most k instructions to be reordered, provided that one is write and another

is a read instruction. This *reorder-bounded analysis* was also shown effective by [18] and seems a natural way to make this problem tractable.

4 Unbounded Buffer Analysis

In this section we show that for any finite data domain program and safety property ϕ , there exists a buffer size k_0 such that it is sufficient to prove ϕ for all buffers up-to size k_0 . Note that for programs with write instructions inside loops, it is possible to keep on writing to the buffer without flushing them to the main memory. However since the data domain is finite, such instruction are guaranteed to repeatedly write the *same set of values* to the buffer. It is this repetition that guarantees the existence of a sufficient bound on the buffer.

The set of states in TSO_k are monotonic with respect to the buffer bound, i.e. $S_k \subseteq S_{k'}$, for all $k \leq k'$. Let $s_{\downarrow (cs, Gm, Lm, Buff_{lst})}$ denote the restriction of a state in S to only control states, global memory, local memory, and last writes (if any) to shared variables in buffers. Let $S_k \downarrow_{(cs, Gm, Lm, Buff_{lst})} = \{s_{\downarrow (cs, Gm, Lm, Buff_{lst})} \mid s \in S_k\}$ be the states of S_k after projecting out the above information. For finite data domain the set $\bigcup_{k=0}^{\infty} S_k \downarrow_{(cs, Gm, Lm, Buff_{lst})}$ is finite because only finitely many different possibilities exist for functions cs , Gm , Lm and $Buff_{lst}$. Further, $S_k \downarrow_{(cs, Gm, Lm, Buff_{lst})} \subseteq S_{k+1} \downarrow_{(cs, Gm, Lm, Buff_{lst})}$. Therefore there exists a k_0 such that $S_{k_0} \downarrow_{(cs, Gm, Lm, Buff_{lst})}$ is equal to the set $S_{k_0+1} \downarrow_{(cs, Gm, Lm, Buff_{lst})}$. In this section we show that for every $k > k_0$, sets $S_k \downarrow_{(cs, Gm, Lm, Buff_{lst})}$ and $S_{k+1} \downarrow_{(cs, Gm, Lm, Buff_{lst})}$ are equal and hence we can stop the analysis at k_0 .

For a buffer $\text{Buff}(t)$, let $\sigma_{\text{Buff}(t).lst}$ denote the sequence of last writes to shared variables in buffer $\text{Buff}(t)$. Let $\text{Exec}(Gm, Lm(t), \text{Buff}(t), \sigma.a.\sigma', (x, v), Lm'(t))$ be a predicate, where $a = (x := e)$, that holds true iff (i) after executing sequence $\sigma_{\text{Buff}(t).lst}.\sigma.a.\sigma'$ from the global memory Gm and local memory $Lm(t)$ the local memory of process t is $Lm'(t)$ and (ii) in the same sequence the value of expression e in write instruction $x := e$ at label a is v . The following two lemmas relate the states of TSO_k and TSO_{k+1} transition systems. We use $s_0 \xrightarrow{\sigma} s$ to denote a sequence of transitions over a sequence σ of labels.

Lemma 1. *For all $n, \sigma, s \in S_{k+1}$ such that $s_0 \xrightarrow{\sigma} s$, $|\sigma| = n$ and $k \geq 0$, there exists a state $s' \in S_k$ such that $s'.Gm = s.Gm$ and for all $t \in \text{TID}$,*

1. $(|s.\text{Buff}_{k+1}(t)| = k + 1) \Rightarrow \exists x, v.$

$$\left\{ \begin{array}{l} (i) \ s.\text{Buff}_{k+1}(t) = s'.\text{Buff}_k(t).(x, v) \\ (ii) \ \exists q_t, q_{t'}, \sigma', \sigma'' \text{ such that } (q_t, a, q_{t'}) \in \delta_t, \\ \quad s'.cs(t) \xrightarrow{\sigma'} q_t, q_{t'}.cs(t) \xrightarrow{\sigma''} s.cs(t), \\ \quad \sigma' \text{ and } \sigma'' \text{ do not modify the buffer of } P_t, \text{ and} \\ \quad \text{Exec}(s'.Gm, s'.Lm(t), s'.\text{Buff}_k(t), \sigma'.a.\sigma'', (x, v), s.Lm(t)) \end{array} \right.$$

and
2. $(|s.\text{Buff}_{k+1}(t)| < k + 1) \Rightarrow$

$$\left\{ \begin{array}{l} (i) \ s.\text{Buff}_{k+1}(t) = s'.\text{Buff}_k(t), \\ (ii) \ s.Lm(t) = s'.Lm(t), \text{ and} \\ (iii) \ s.cs(t) = s'.cs(t) \end{array} \right.$$

$s'.cs(t) = s.cs(t)$. The above lemma states that every state in S_{k+1} where the buffer sizes of all processes are less than $k+1$, is also present in S_k . The detailed proof of this lemma is given in the Appendix. Now we are ready to prove that after k_0 , any increase in buffer size does not yield any new reachable control location.

Theorem 1. *For all k ,*

$$\begin{aligned} (S_k \downarrow_{(cs, Gm, Lm, Buff_{lst})} = S_{k+1} \downarrow_{(cs, Gm, Lm, Buff_{lst})}) &\Rightarrow \\ (S_{k+1} \downarrow_{(cs, Gm, Lm, Buff_{lst})} = S_{k+2} \downarrow_{(cs, Gm, Lm, Buff_{lst})}) &\end{aligned}$$

Proof. there exists a state $s' \in S_{k+1}$ such that $s.cs = s'.cs$, $s.Gm = s'.Gm$, $s.Lm = s'.Lm$ and $s.Buff_{lst} = s'.Buff_{lst}$. It is sufficient to show that $(S_{k+2} \downarrow_{(cs, Gm, Lm, Buff_{lst})} \subseteq S_{k+1} \downarrow_{(cs, Gm, Lm, Buff_{lst})})$ as the other side of inclusion holds. Let us prove it by contradiction, i.e. there is a state $s \in S_{k+2}$ such that no state $s' \in S_{k+1}$ exists with $s.cs = s'.cs$, $s.Gm = s'.Gm$, $s.Lm = s'.Lm$ and $s.Buff_{lst} = s'.Buff_{lst}$. Following Lemma 1, this state s must have at least one buffer with $k+2$ entries in it. Without loss of generality let $t \in TID$ such that $s.Buff_{k+2}(t)$ is the only full buffer.

1. Clearly, there exists a state $s' \in S_{k+2}$ where all buffers except t are the same as in s , $s'.Buff_{k+2}(t)$ is of size $k+1$ and there exists a sequence of transitions $\sigma.a.\sigma'$ from $s'.cs(t)$ to $s.cs(t)$ by process t with only one write operation a .
2. For state s' , the conditions $s'.Gm = s.Gm$ (as no flush operation in σ), and $Exec(s'.Gm, s'.Lm(t), s'.Buff_{lst}(t), \sigma.a.\sigma', s.Lm(t))$ hold.
3. As all buffers of s' are of size at most $k+1$ therefore s' also exists in S_{k+1} (Lemma 1).
4. As $S_k \downarrow_{(cs, Gm, Lm, Buff_{lst})} = S_{k+1} \downarrow_{(cs, Gm, Lm, Buff_{lst})}$ holds, therefore there exists a state $s'' \in S_k$ such that (i) $s''.Gm = s'.Gm$, (ii) $s''.Lm = s'.Lm$, (iii) $s''.cs = s'.cs$, and (iv) $s''.Buff_{lst}(t) = s'.Buff_{lst}(t)$ for all $t \in TID$.
5. This state s'' can have at most k entries in its process buffers. Therefore this state must be present in S_{k+1} as well.
6. Using Point 2 and the conditions (i),(ii),(iii), and (iv) of Point 4 above, we get $Exec(s''.Gm, s''.Lm(t), s''.Buff_{lst}(t), \sigma.a.\sigma', s.Lm(t))$. This implies that after executing the sequence $\sigma.a.\sigma'$ by process t from state s'' in S_{k+1} the resultant state, say s''' will have at most $k+1$ write entries in the buffer of process t . Further the global memory, local memories, control states and last writes to shared variables in buffers will be identical in s''' and s . Therefore $s''' \in S_{k+1}$ is the matching state with respect to s , a contradiction.

5 Trace partitioning approach

As a consequence of Theorem 1 one can use an explicit state model checker for state reachability analysis of finite data domain programs. However, in this paper we are interested in adapting a recently proposed trace partitioning based verification method [16,25] for relaxed memory models. This method has been

shown very effective for verification under the SC memory model. The approach for SC verification, as given in [25], is presented in Algorithm 1. Firstly, an automaton is built that represents the set of symbolic traces under the SC memory model. For SC memory model such an automaton is obtained by language level shuffle operation [26,17] on individual processes. Subsequently, a symbolic trace is picked from this automaton and checked against a given safety property using weakest precondition axioms [15]. If this trace violates the given property then we have a concrete erroneous trace. Otherwise, an alternating finite automaton (AFA) [12] is constructed from the proof of correctness of this trace.

The AFA construction algorithm ensures that every trace in the language of this AFA is correct and hence can be safely removed from the set of all symbolic traces of the input program. This process is repeated until either all symbolic traces are proved correct or an erroneous trace is found. This algorithm is sound and complete for finite data domain programs.

Input: A concurrent program $\mathcal{P} = \{p_1, \dots, p_n\}$ with safety property ϕ
Result: *yes*, if program is safe else a counterexample
Construct the automaton $\mathcal{A}(\mathcal{P})$ to capture the set of all SC traces of \mathcal{P} ;
Let **tmp** be the language of $\mathcal{A}(\mathcal{P})$;
while **tmp** *is not empty* **do**
 Let $\sigma \in \mathbf{tmp}$ with ϕ as a safety assertion to be checked;
 Let $\hat{\mathcal{A}}_{\sigma, \neg\phi}$ be the AFA constructed from σ and $\neg\phi$;
 if σ *violates* ϕ **then**
 σ is a valid counterexample;
 return (σ);
 else
 tmp := **tmp** \setminus *Rev*, where *Rev* is the reverse of the language of $\hat{\mathcal{A}}_{\sigma, \neg\phi}$;
 end
end
return (*yes*);

Algorithm 1: SC verification algorithm[25]

The main challenge in applying this trace partitioning approach to the TSO memory model is the construction of the set of symbolic traces. Consider a program with two processes in Figure 3. With initial values of shared variables x and y as 0, it is possible to have $\ell_1 = \ell_2 = 0$ under the TSO memory model. We can construct a symbolic trace **b.d.a.c** such that after executing this sequence the state $\ell_1 = \ell_2 = 0$ is reached.

a. $x:=1$ | c. $y:=1$
b. $\ell_1:=y$ | d. $\ell_2:=x$

Fig. 3

Note that this trace is not constructible using the standard interleaving semantics which was used to construct the set of traces under the SC memory model. This is because of the program order between **a** and **b** in process 1 and between **c** and **d** in process 2. To use Algorithm 1 for the TSO memory model we

would like to first construct a set of all such symbolic traces such that the sequential executions of these traces yield all reachable states under the TSO memory model. For the above example, it involves breaking the program orders **a** – **b** and **c** – **d** and then applying standard interleaving semantics to construct symbolic

traces under the TSO memory model. Let us look at another non-trivial example in Figure 4.

Assume the initial values of all variables are 0, and ℓ, m are local variables. In TSO it is possible to have the final values of variables ℓ and m as 0. This can happen when writes at **b** and **e** are still in the buffers and the read operations at **c** and **f** read from the initial values. Let us construct a symbolic trace whose sequential execution will yield this state. In this trace label **e** must appear after label **c** and label **b** must appear after label **f**. This means that the trace will break either the order between **b** and **c** or the order between **e** and **f**. However, by breaking the order between **b** and **c** the value of $\ell = 2$, assigned at **a**, no longer flows to **b** and hence y is assigned the wrong value 1. Similarly by breaking the order **e** and **f** the value of $m = 3$, assigned at **d** no longer flows to **e** and hence x is assigned the wrong value 1. In a nutshell, it is not possible to create a symbolic trace whose execution will yield the state where $\ell = m = 0$, $x = 5$, and $y = 3$. Notice that the problem appeared because of the use of the same local variable in two definitions. Such a scenario is unavoidable when (i) multiple reads are assigned to the same local variable, and/or (ii) in the case of loops the local variable appears in a write instruction within the loop.

We propose to handle such cases by renaming local variables, viz. ℓ and m in this case. For example, the execution of trace $\sigma = \ell := 2. \ell_1 := \ell. \ell := x. y := \ell_1 + 1. m := 3. m_1 := m. m := y. x := m_1 + 2$ results in state $\ell = m = 0, x = 5, y = 3$ as required by a TSO execution. Let us look at instructions highlighted in gray color more carefully. We earlier saw that the problem arises when reordering **b** – **c** and **e** – **f** instructions as their reordering will break the value flows of ℓ and m from **a** and **d** respectively. Therefore, we create new instances of these local variables, ℓ_1 and m_1 , to take the *snapshot* of ℓ and m respectively which are later used in the write instructions **b** and **e**. This renaming ensures that even if we reorder **b** – **c** and **e** – **f** instructions (as done in σ) the correct value flows from **a** to **b** and from **d** to **e** are not broken. We will show that for a buffer bound of k it is sufficient to use at most k instances of these local variables and they can be safely reused even in the case of loops. We call such symbolic traces, that correspond to TSO_k executions, as *SC interpretable traces*. Formally, *SC interpretation* of a trace $\sigma \in \text{LABL}^*$ is a function $\text{SCI} : \text{LABL}^* \times \text{Var} \rightarrow \text{Val} \cup \{\text{Undef}\}$, such that $\text{SCI}(\sigma, x)$ calculates the last value assigned to variable x in the sequential execution of σ . For example, if $\sigma = \mathbf{a.b.c}$ where labels **a**, **b**, and **c** denote $\ell := 3, x := \ell + 2$ and $y := 2$ respectively then $\text{SCI}(\sigma, x) = 5$ and $\text{SCI}(\sigma, \ell) = 3$. Label **Undef** is used to denote the in-feasibility of σ as some boolean expressions in **assume** instructions may become unsatisfiable because of the values that flow in them. If σ does not contain any assignment to x then $\text{SCI}(\sigma, x)$ returns the initial value of x .

Let us now construct a transition system such that the traces of this transition system represent SC interpretable traces corresponding to TSO_k semantics. We represent this transition system as $\text{TSO}_k^\# = \langle S^\#, \Rightarrow_k, s_0^\# \rangle$. Every state $s^\# \in S^\#$ is of the form $(cs, Li, \text{Buff}_k^\#)$ such that $cs : \text{TID} \rightarrow Q$ represents process control

states, and $\text{Buff}_k^\# : \text{TID} \rightarrow (\text{SV} \times \text{LABL})^k$ represents per process buffers of length k . Unlike the buffers of TSO_k , these buffers contain write instruction labels along with the modified shared variable. A function $\text{Li} : \text{TID} \times \text{LV} \rightarrow \mathbb{N}$ tracks the instances of the local variables which have been used (for renaming purposes) in the construction of traces up to a given state. First, we define \Rightarrow_k for simple

$$\begin{array}{c}
\frac{\text{Ins}(a) = (\ell := x), \quad \text{Buff}_k^\#(t) \downarrow_{\{x\} \times \text{LABL}} = \epsilon}{(\text{cs}, \text{Li}, \text{Buff}_k^\#) \xRightarrow{a}_k (\text{cs}', \text{Li}, \text{Buff}_k^\#)} \quad (\text{MREAD}^\#) \\
\\
\frac{\text{Ins}(a) = (\ell := e)}{(\text{cs}, \text{Li}, \text{Buff}_k^\#) \xRightarrow{a}_k (\text{cs}', \text{Li}, \text{Buff}_k^\#)} \quad (\text{LWRITE}^\#) \\
\\
\frac{\text{Ins}(a) = (\text{assume}(e))}{(\text{cs}, \text{Li}, \text{Buff}_k^\#) \xRightarrow{a}_k (\text{cs}', \text{Li}, \text{Buff}_k^\#)} \quad (\text{ASSUME}^\#) \\
\\
\frac{\text{Buff}_k^\# = (x, a). \text{Buff}_k^{\#'}}{(\text{cs}, \text{Li}, \text{Buff}_k^\#) \xRightarrow{a}_k (\text{cs}', \text{Li}, \text{Buff}_k^{\#'})} \quad (\text{FLUSH}^\#)
\end{array}$$

Fig. 5: All rules assume transitions for thread t , ie. $\text{cs}[t] = q$, $(q, a, q') \in \delta_t$, and $\text{cs}' = \text{cs}[t \leftarrow q']$

cases, viz. read from memory, operations associated with local variables like $\text{assume}(e)$ and $\ell := e$, and non-deterministic flush. In Rules $\text{MREAD}^\#$, $\text{LWRITE}^\#$, and $\text{ASSUME}^\#$ the labels that denote these operations are put in the trace with only change in the control state of the process. As there is no notion of local and global valuation in a state $\mathbf{s}^\#$ of the transition system, no update takes place unlike in TSO_k . For memory read operation, in Rule $\text{MREAD}^\#$, the condition on the buffer of P_t is the same as in TSO_k . For non-deterministic flush operation, Rule $\text{FLUSH}^\#$ removes the first label present in the buffer of P_t and puts that in the trace. In rule $\text{ASSUME}^\#$, the assume instruction is simply put in the trace without evaluating the satisfiability of the boolean expression. This is different from the corresponding rule in TSO_k . This difference follows from the fact that we are only interested in constructing symbolic traces. Symbolic model checking of these traces will ensure that only feasible executions get analyzed (where all assume instructions hold true). Now let us look at the remaining three operations, viz. read from the buffer, write to the buffer and fence instruction, in detail.

Buffered Read Like TSO_k , this transition takes place when P_t executes an instruction $\ell := x$ to read the value of shared variable x and store it in its local variable ℓ .

$$\frac{\begin{array}{l} \text{Ins}(a) = (\ell := x), \text{Buff}_k^\# \downarrow_{\{x\} \times \text{LABL}} = \alpha.(x, b), \\ \text{Ins}(b) = (x := e), \text{Ins}(c) = (\ell := e) \end{array}}{(\text{cs}, \text{Li}, \text{Buff}_k^\#) \xRightarrow{c}_k (\text{cs}', \text{Li}, \text{Buff}_k^\#)} \quad (\text{BREAD}^\#)$$

For this transition to take place, the buffer of P_t must have at least one write instruction that modifies the shared variable x . Conditions $\text{Buff}_k^\# \downarrow_{\{x\} \times \text{LABL}} = \alpha.(x, b)$ and $\text{Ins}(b) = (x := e)$ ensure that the last write to x in $\text{Buff}_k^\#$ of P_t is due to instruction $\text{Ins}(b)$ which is of the form $(x := e)$. Under these conditions, in TSO_k , read of x uses the value of expression e to modify ℓ . Whereas in $\text{TSO}_k^\#$ a label c is added to the trace such that $\text{Ins}(c)$ represents the assignment of e to variable ℓ .

Buffered Write This transition takes place when P_t executes a write instruction of the form $x := e$. Let $\vec{\ell}$ be a set of local variables used in expression

e. For each of the local variables ℓ in $\vec{\ell}$, an integer $\text{Li}(\ell)$ is used to create an assignment instruction of the form $\ell_{\text{Li}(\ell)} := \ell$. These instructions are put in the trace (through corresponding symbolic labels $\vec{a}_{\vec{\ell}}$). Further, expression e is also modified where every instance of a local variable ℓ in $\vec{\ell}$ is substituted with $\ell_{\text{Li}(\ell)}$.

$$\begin{array}{c}
\text{Ins}(a) = (sv := e), \text{FV}(e) = \vec{\ell}, |\text{Buff}^\sharp_k| < k, \\
\forall \ell \in \vec{\ell}, \text{ create a label } a_\ell \text{ (if not already present in LABL)} s.t. \\
\text{Ins}(a_\ell) = (\ell_{\text{Li}(\ell)} := \ell), \text{Li}'[\ell] = \text{Li}[\ell] \% (k+1) + 1 \\
\text{ create a label } a' \text{ (if not already present in LABL)} s.t. \\
\text{Ins}(a') = (sv := e'), e' = e[\vec{\ell} / \vec{\ell}_{\text{Li}(\ell)}], \\
\text{Buff}^\sharp_k = \text{Buff}^\sharp_k[t \leftarrow \text{Buff}^\sharp_k[t].(sv, a')] \\
\hline
(cs, \text{Li}, \text{Buff}^\sharp_k) \xrightarrow{\vec{a}_{\vec{\ell}}} (cs', \text{Li}', \text{Buff}^\sharp_k) \quad (\text{BWRITE}^\sharp)
\end{array}$$

This modified expression e' is denoted $e[\vec{\ell} / \vec{\ell}_{\text{Li}(\ell)}]$ in Rule BWRITE^\sharp . A label, a' , representing the assignment of e' to x is put in the buffer in the form of a tuple (x, a') . Note that the transition rule BWRITE^\sharp increases the value of $\text{Li}(\ell)$ (modulo $(k+1)$) for every local variable ℓ present in expression e . We can show the following property,

Lemma 2. *For a state $s^\sharp = (cs, \text{Li}, \text{Buff}^\sharp_k)$ of TSO^\sharp_k , if $\text{Li}(\ell) = m$ then local variable ℓ_m does not appear in any write instruction used in buffers Buff^\sharp_k .*

Proof. Suppose $\text{Li}(\ell) = m$ holds. By assumption, local variables among processes are disjoint therefore the only possibility is that $\text{Buff}^\sharp_k[t]$ contains a write instruction that uses local variable ℓ_m . If this were the case then there must be at least $k+1$ different writes appearing between that write and the time s^\sharp is reached. This holds because every write, that uses a local variable ℓ first increments its index by 1 and wraps around after $k+1$. This incremented index is then used to create an instance of the local variable ℓ used in this write operation. But it contradicts our assumption that the buffer is of bounded length k .

The above lemma is used in the equivalence proof of TSO_k and TSO^\sharp_k .

Fence Fence instruction, like TSO_k , gets enabled only when $\text{Buff}^\sharp_k[t]$ is empty. In the resultant state, function $\text{Li}(t, \ell)$ is set to 1 for every local variable ℓ of Process P_t . This enables the reuse of indices in Function Li while preserving Lemma 2.

$$\begin{array}{c}
\text{Ins}(a) = (\text{fence}), \text{Buff}^\sharp_k[t] = \epsilon \\
\text{Li}' = \text{Li}[(t, \ell) \leftarrow 1], \forall \ell \in \text{LV}_t \\
\hline
(cs, \text{Li}, \text{Buff}^\sharp_k) \xrightarrow{\epsilon} (cs', \text{Li}', \text{Buff}^\sharp_k) \quad (\text{FENCE}^\sharp)
\end{array}$$

To show the equivalence of TSO_k and TSO^\sharp_k we want to prove the following; (i) for every state s reachable in TSO_k there exists a trace σ^\sharp in TSO^\sharp_k such that the *SC interpretation* of σ^\sharp reaches a state with the same global memory and local memory as of s , and (ii) for every trace σ^\sharp of TSO^\sharp_k such that its *SC interpretation* is not **Undef** (i.e. execution should be feasible) there exists a state $s \in \text{TSO}_k$ with same global and local memory as obtained after the *SC interpretation* of σ^\sharp . We formally prove the following theorem in the Appendix.

Theorem 2. *Transition systems TSO_k and $\text{TSO}_k^\#$ are equivalent in terms of state reachability.*

In Theorem 1 we used the restricted set $S_{\downarrow \text{cs}, \text{Gm}, \text{Lm}, \text{Buff}_{l, st}}$ as a means to define fixed point. However, there are no explicit representations of the global memory (Gm) and the local memory (Lm) in the state definition of $\text{TSO}_k^\#$. Therefore, in order to define a fixed point condition like Theorem 1 we first augment the definition of state in $\text{TSO}_k^\#$ to include global memory and local memory. Let $\text{Gm}^\# : \text{SV} \rightarrow \text{Lab}$ and $\text{Lm}^\# : \text{TID} \times \text{LV} \rightarrow \text{Lab}$ be the functions assigning labels (of write instructions) to shared variables and local variables respectively. Specifically, $\text{Gm}^\#(x) = a$ means that the write instruction at label a was used to define the current value of x in this state. Similarly, $\text{Lm}^\#(t, \ell) = a$ means that the write instruction at label a was used to define the current value of local variable ℓ of process t . Note that in the construction of $\text{TSO}_k^\#$ the values written by these write instructions are only being represented symbolically using instruction labels. Therefore we need a way to relate the instruction labels and the actual values written. For a concurrent program P with finite data domain it is possible to construct an equivalent program P' such that every assignment to variables in P' is only of constant values. For example, consider the program in Figure 6 such that the domain of variable x is $\{1, 2\}$. This program is equivalent to the program in Figure 7 where only constant values are used in the write instructions. Here the domain of x is used along with if-then-else conditions to decide the value that needs to be written to y .

		After this transformation, every write label uniquely identifies the value written to a shared variable. Hence the functions $\text{Gm}^\#, \text{Lm}^\#$ can be extended to $\text{SV} \rightarrow \text{Val}$ and $\text{LV} \rightarrow \text{Val}$ respectively. This allows us to use Theorem 1 for checking the fixed point.
	$\ell := x$	
	if ($\ell = 1$)	
	$y := 4$	
	else if ($\ell = 2$)	
$\ell := x$	$y := 5$	
$y := \ell + 3$		5.1 Fence Insertion For Program Correction

Fig. 6

Fig. 7

Let P be a program that is correct under the SC memory model. Let σ be an execution of P that violates the given safety property under the TSO memory model. We can insert a fence instruction in P so that σ does not appear as an execution under the TSO memory model. Towards this we use the *critical cycle* based approach of [27] and [6] to detect the locations of fence insertions. For an execution σ , let Cmpt_σ be a *competing*[6] or *conflicting*[27] relation on the read and write events of σ such that $(a, b) \in \text{Cmpt}_\sigma$ iff (i) both memory events operate on the same location but originate from different processes, (ii) at least one of them is a write instruction, and (iii) a appears before b in σ . Let po_σ denote the program order among instructions of processes present in σ . This is defined based on the process specification. Let $\text{ppo}_\sigma = \text{po}_\sigma \setminus \{(a, b) \mid a \in W, b \in R, (a, b) \in \text{po}_\sigma\}$ be a subset of po_σ preserved under TSO memory model, i.e. everything except write-read orders.

An Execution σ contains a critical cycle $\xrightarrow{cs} \subseteq (\text{Cmpt}_\sigma \cup \text{ppo}_\sigma)^+$ iff (i) no cycle exists in $(\text{Cmpt}_\sigma \cup \text{ppo}_\sigma)^+$, (ii) per process there are at most two memory accesses

a and b in \xrightarrow{cs} such that $\text{Loc}(a) \neq \text{Loc}(b)$, and (ii) for a given shared variable x there are at most three memory accesses on x which must originate from different processes. Following Theorem 1 of [6], an execution in TSO is sequentially consistent if and only if it does not contain any critical cycle. Therefore, in order to forbid an execution in TSO that is not sequentially consistent, it is sufficient to ensure that no critical cycle exists in that execution. To avoid critical cycles, we need to strengthen the ppo_σ relation by adding a minimal set of program orders such that Point (i) of critical cycle definition is not satisfied, i.e. finding a set $\text{Dlay} \subseteq \text{po}_\sigma \setminus \text{ppo}_\sigma$, set of write-read pairs of instructions within each process, such that $(\text{Cmpt}_\sigma \cup \text{ppo}_\sigma \cup \text{Dlay})^+$ becomes cyclic. Once we identify that minimal set of program orders we insert fence instructions in between them to enforce the required orderings.

Overall Algorithm Algorithm that combines incremental buffer bounded verification and fence insertion for finite data programs works as follows. We start the verification with buffer bound of 0. Towards this, the transition system $\text{TSO}^\#_k$ is constructed using the relation \Rightarrow_k given in this section. This transition system is represented as an automaton with error location representing the accepting states and initial locations representing the initial state. The set of traces accepted by this automaton are the passed to the trace partitioning algorithm implemented by [25] in the tool **ProofTraPar**. If an erroneous trace is found then the program is not safe even under the SC memory model and hence the algorithm returns the result as ‘Unsafe’. If all traces satisfy the given safety property then the bound is increased by one and the analysis starts again. If an error trace is found for non-zero buffer bound then the critical cycles are obtained from this trace. Using these critical cycles a set of fence locations are generated and the input program is modified by inserting fences in the code. After the modification the analysis again starts with the same bound. This is just an implementation choice because even if we increase the bound after the modification still the fixed point will be eventually reached.

6 Experimental Results

We implemented our approach by extending the tool **ProofTraPar** which implements the trace partitioning based approach of [25]. We implemented $\text{TSO}^\#_k$ semantics and fixed point reachability check on top of **ProofTraPar**. Its performance was compared against **memorax** which implements sound and complete verification of state reachability under the TSO memory model. Note that other tools which exist in this landscape of relaxed memory verification either consider SC behaviour as specification [3,6,10] or are sound but not complete [23,1,28]. However **memorax** does not assume any bound on the buffer size and it uses the coverability based approach of well-structured-transition systems. Table 8 compares the performance, in terms of time and memory, of our approach with **memorax**. We ran all experiments on Intel i7-3.1GHz, 4 core machine with 8GB RAM. Out of 11 examples, our tool outperformed **memorax** in 8 examples. Our

Program	# P	ProofTraPar		Memorax[2]		# F
		Time (Sec)	Memory (MB)	Time (Sec)	Memory (MB)	
Peterson.safe	2	1.19	20	0.9	43	2
Dekker.safe	2	1.6	21.3	54.2	676	2
Lamport.safe	2	17	42	97	2312	4
Szymanski.safe	2	27	121	ERR	ERR	4
Alternating Bit(ABP)	2	3.12	39	0.17	11	0
Dijkstra	2	16	70	-	-	2
Pgsql	2	1.2	20	210	2800	2
RWLock.safe (2R,1W)	3	41	164	-	-	2
clh	2	326	1500	-	-	0
Simple-dekker	3	103	155	600	3280	3
Qrcu.safe (2R,1W)	3	490	3000	-	-	0

Fig. 8: Comparison of our tool with Memorax[2]. Time out, denote by ‘-’ is set to 10 minute. #P and #F denote number of processes and number of fences synthesized.

tool not only performed better in terms of time but also in terms of the memory consumption. Except in two cases, *qrcu* and *clh queue*, on every other example our tool consumed less than 200 MB of RAM. Whereas **memorax** in most cases took more than 500 MB of RAM and in some cases even touched the 3GB mark. Programs like *Alternating bit protocol*, *clh queue* and *Qrcu* (quick read copy update algorithm) remain correct even under TSO memory model. For other algorithms where bugs were exposed under TSO we were able to synthesize fences to correct their behaviour.

Analysis of the benchmarks **memorax** performed better on three benchmarks, viz. *peterson*, *szymanski*, and *ABP*. After carefully looking at them we realized that the performance of **memorax** loosely depends upon the number of backward control flow paths from error location to the start location, and number of write instructions present along those paths. In benchmarks where **ProofTraPar** outperformed **memorax**, viz. *dekker*, *lamport*, *clh*, *qrcu*, more than two such control paths exist. To further check this hypothesis experimentally we modified *peterson* and *ABP* to add a write instruction along an already existing control flow path where no write instruction was present. This write was performed on a variable which was never read and hence did not affect the program. After this modification **memorax** became more than 6 time slower in analyzing these two benchmarks. Further, the analysis of these modified benchmarks with **ProofTraPar** exhibited a very little (less than a second) increase in time as compared to the unmodified benchmarks. Interestingly, a bug was exposed in **memorax** when we made a similar change in *szymanski*. As a result of this bug the modified program *szymanski* was declared as safe. Note that the original program *szymanski* is incorrect under TSO and we only modified the code by adding a write instruction to an unused variable. Therefore it is not possible for

the modified program *szymanski* to become safe unless there is a bug in the tool. This bug was also confirmed by the author of *memorax*.

6.1 Discussion

Note that *memorax* starts from the symbolic representation of all possible configurations of buffer contents which it further refines using backward reachability analysis. However, in our approach we start from a finite and small buffer bound (an under-approximation) and keep expanding until we reach a fix point. We believe that this difference, picking an over-approximation as a starting point in one case and an under-approximation as a starting point in the other case, plays a crucial role in the better performance of our approach on these benchmarks.

In all benchmarks, except *peterson*, buffer size of 1 was sufficient to expose the error. In *peterson*, buffer size of 2 was needed to expose the bug. Effectively, the buffer size depends upon the minimum distance (along control flow path) between a write and a read instruction within a process whose reordering reveals the bug. In the case of *peterson*, this distance is 2 since the reordering of first instruction (write to *flag_i*) and third instruction (read of *flag_j*) within each process reveals the bug. In our benchmarks fence instructions were inserted after finding an erroneous trace, as discussed in Section 5.1. Fence instruction restricts the unbounded growth of the buffer by flushing the buffer contents. As a result, when a fence is inserted within a loop the buffer never grows in size with loop iterations and fix point is reached quickly. In fact, for all the benchmarks, if a bug was exposed with buffer size k then after inserting the fence instruction the fix point was reached with buffer size $k + 1$. Benchmarks which remain correct under TSO, a larger buffer bound was required to reach the fix point and this bound depends upon the number of write operations in each process. As a result, their analysis took longer time and consumed more memory. Detailed analysis of the benchmarks and the tool are available at www.cse.iitd.ac.in/~chinmay/ProofTraParTSO.

7 Conclusion and Future Work

This paper uses the trace partitioning based approach to verify state reachability of concurrent programs under the TSO memory model. We have also shown that for finite state programs there exists a buffer bound such that if program is safe up-to that bound then the program is guaranteed to be safe for unbounded buffers as well. This work can be easily extended to PSO memory model as well. This method gives us an alternate decidability proof of state reachability under TSO (and PSO) memory model. We have also shown experimentally that for standard benchmarks used in the literature such a bound is very small (in the range of 2-4) and hence we may use SC verification based methods to efficiently check concurrent programs under these memory models. We believe that for other buffer based memory models a buffer bound can be shown to exist in a similar manner. Recently [21] proposed a buffer based operational semantics for C11 model. It will be interesting to investigate the use of bounded buffer based method proposed in this paper to that semantics as well.

References

1. ABDULLA, P. A., ARONIS, S., ATIG, M. F., JONSSON, B., LEONARDSSON, C., AND SAGONAS, K. F. Stateless model checking for TSO and PSO. In *TACAS'15*.
2. ABDULLA, P. A., ATIG, M. F., CHEN, Y.-F., LEONARDSSON, C., AND REZINE, A. Counter-example guided fence insertion under tso. *TACAS'12*, Springer-Verlag.
3. ABDULLA, P. A., ATIG, M. F., AND NGO, T.-P. The best of both worlds: Trading efficiency and optimality in fence insertion for tso. In *ESOP'15*, Springer-Verlag.
4. ALGLAVE, J., KROENING, D., NIMAL, V., AND POETZL, D. Don't sit on the fence - A static analysis approach to automatic fence insertion. In *CAV'14*.
5. ALGLAVE, J., KROENING, D., NIMAL, V., AND TAUTSCHNIG, M. Software verification for weak memory via program transformation. In *ESOP'13* (2013).
6. ALGLAVE, J., AND MARANGET, L. Stability in weak memory models. *CAV'11*, Springer-Verlag, pp. 50–66.
7. ATIG, M. F., BOUAIJANI, A., BURCKHARDT, S., AND MUSUVATHI, M. On the verification problem for weak memory models. *SIGPLAN Not., Jan'10* 45, 1.
8. ATIG, M. F., BOUAIJANI, A., AND PARLATO, G. Getting rid of store-buffers in TSO analysis. In *CAV'11* (2011).
9. BOUAIJANI, A., CALIN, G., DEREVENETC, E., AND MEYER, R. Lazy TSO reachability. In *FASE'15* (2015).
10. BOUAIJANI, A., DEREVENETC, E., AND MEYER, R. Checking and enforcing robustness against tso. *ESOP'13*, Springer-Verlag, pp. 533–553.
11. BURNIM, J., SEN, K., AND STERGIOU, C. Sound and complete monitoring of sequential consistency for relaxed memory models. *TACAS'11*, Springer-Verlag.
12. CHANDRA, A. K., KOZEN, D. C., AND STOCKMEYER, L. J. Alternation. *J. ACM* 28, 1 (Jan. 1981), 114–133.
13. DAN, A. M., MESHMAN, Y., VECHEV, M. T., AND YAHAV, E. Predicate abstraction for relaxed memory models. *SAS'13*, pp. 84–104.
14. DAN, A. M., MESHMAN, Y., VECHEV, M. T., AND YAHAV, E. Effective abstractions for verification under relaxed memory models. *VMCAI'15*, pp. 449–466.
15. DIJKSTRA, E. W. Guarded commands, nondeterminacy and formal derivation of programs. *Commun. ACM* 18, 8 (Aug. 1975), 453–457.
16. FARZAN, A., KINCAID, Z., AND PODELSKI, A. Inductive data flow graphs. In *POPL'13*.
17. HOPCROFT, J. E., MOTWANI, R., AND ULLMAN, J. D. Introduction to automata theory, languages, and computation, 2nd edition.
18. JOSHI, S., AND KROENING, D. Property-driven fence insertion using reorder bounded model checking. In *FM 2015: (2015)*.
19. KUPERSTEIN, M., VECHEV, M. T., AND YAHAV, E. Partial-coherence abstractions for relaxed memory models. *PLDI'11*, pp. 187–198.
20. KUPERSTEIN, M., VECHEV, M. T., AND YAHAV, E. Automatic inference of memory fences. *SIGACT News* 43, 2 (2012), 108–123.
21. LAHAV, O., GIANNARAKIS, N., AND VAPEIADIS, V. Taming release-acquire consistency. In *POPL'16* (New York, NY, USA, 2016), POPL 2016, ACM, pp. 649–662.
22. LINDEN, A., AND WOLPER, P. A verification-based approach to memory fence insertion in relaxed memory systems. In *SPIN'11* (2011), pp. 144–160.
23. LINDEN, A., AND WOLPER, P. A verification-based approach to memory fence insertion in pso memory systems. *TACAS'13*, Springer-Verlag, pp. 339–353.
24. MESHMAN, Y., DAN, A. M., VECHEV, M. T., AND YAHAV, E. Synthesis of memory fences via refinement propagation. *SAS'14*, pp. 237–252.

25. NARAYAN, C., SHARMA, S., GUHA, S., AND ARUN-KUMAR, S. From traces to proofs: Proving concurrent program safe (accepted for publishing). *Theoretical Aspects of Software Engineering, 2016* (arXived Version: <http://arxiv.org/abs/1506.07635>).
26. RIDDLE, W. E. An approach to software system modelling and analysis. *Comput. Lang.* 4, 1 (Jan. 1979), 49–66.
27. SHASHA, D., AND SNIR, M. Efficient and correct execution of parallel programs that share memory. *TOPLAS* 10, 2 (Apr. 1988), 282–312.
28. ZHANG, N., KUSANO, M., AND WANG, C. Dynamic partial order reduction for relaxed memory models. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation* (New York, NY, USA, 2015), PLDI 2015, ACM, pp. 250–259.